AQUANTIA®

10G/5G/2.5G/1G/100M Physical Layer PHY

Hot Chips 2015 Conference

Ramin Shirani
Ramin Farjad-Rad

Megatrends Driving Need for Next Generation Connectivity

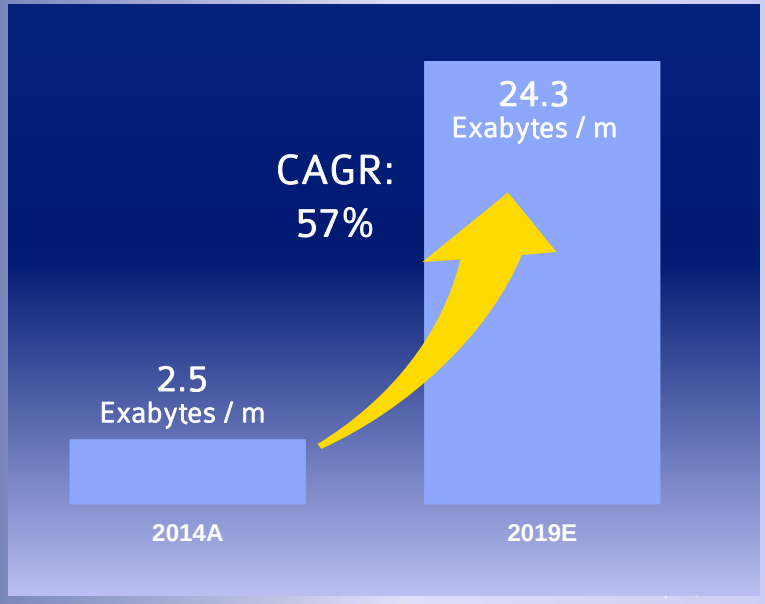

- Transition to Next Generation Ethernet Solution Required to Alleviate Critical Bandwidth Constraints in Global IT Infrastructure

AQUANTIA®

Enterprise Network Structure

- Most enterprise networks follow the Three Tier Design approach
- Approach is proven, widespread, and stable
- More than 90% of enterprise campus wiring = cat5e / cat6
- In 2012, Aquantia developed ICs to increase the speeds on cat5e / cat6 links between enterprise access switches & wireless APs

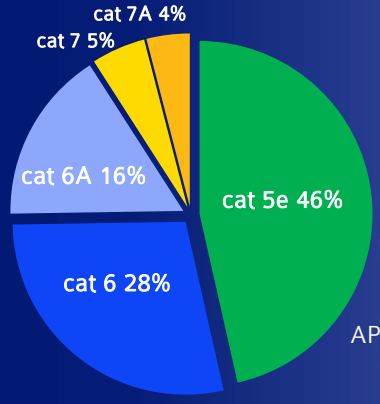

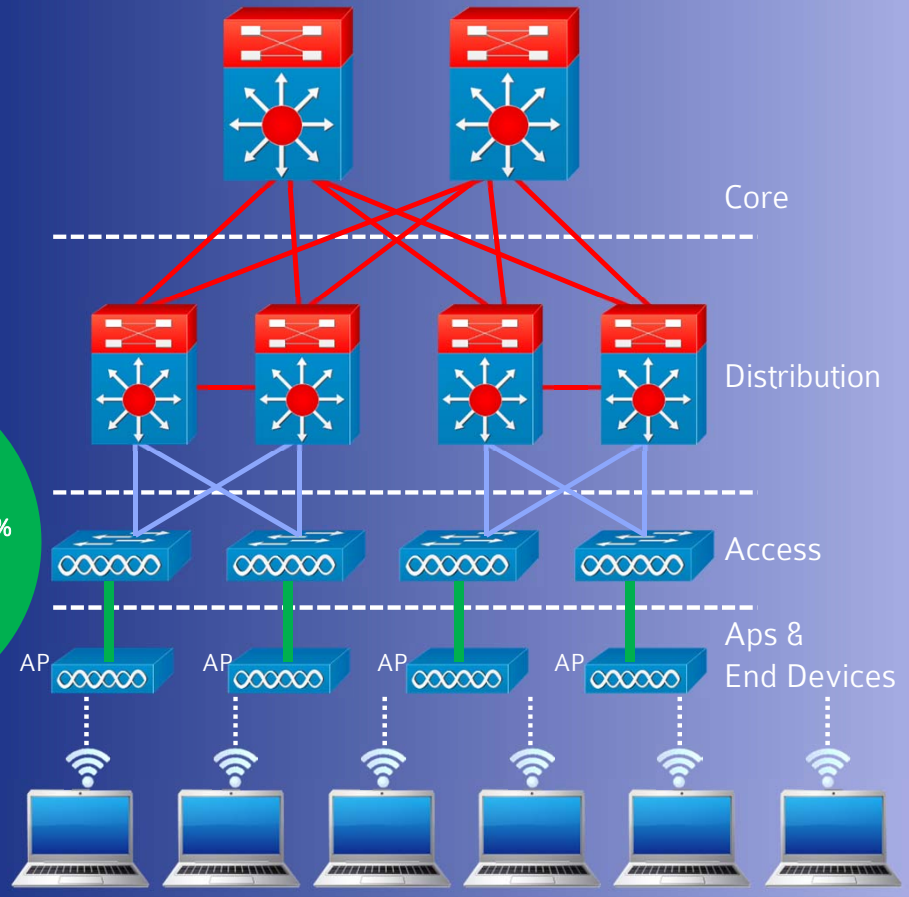

AQUANTIA®

WLAN Upgrade Creating Bottleneck in Enterprise Networks

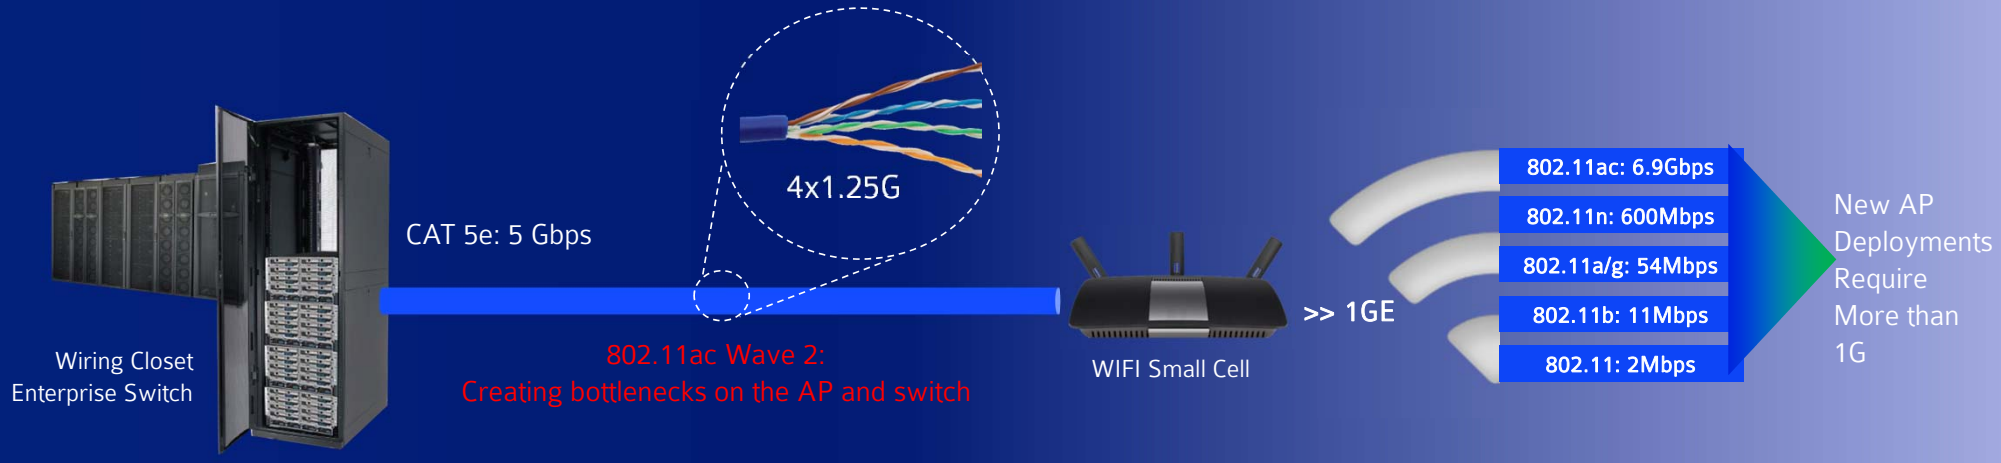

- Legacy wireless access connections are dominated by 1000Base-T with 30W PoE
- Wireless transition to 802.11ac (Wave 1 and Wave 2) needs:
  - Multi-Gigabit speeds in conjunction with 60W UPoE
- Aquantia's 28nm ICs were designed for five speeds on 100 meters:
  - 10G Cat6A and Cat7
  - 5G/2.5G/1G/100M on Cat6A, Cat6, and Cat5e

NBASE-T™ ALLIANCE

Aquantia: founding member

- 2.5G/5G electrical and interoperability specifications

AQUANTIA®

Multi-Gigabit Application Space

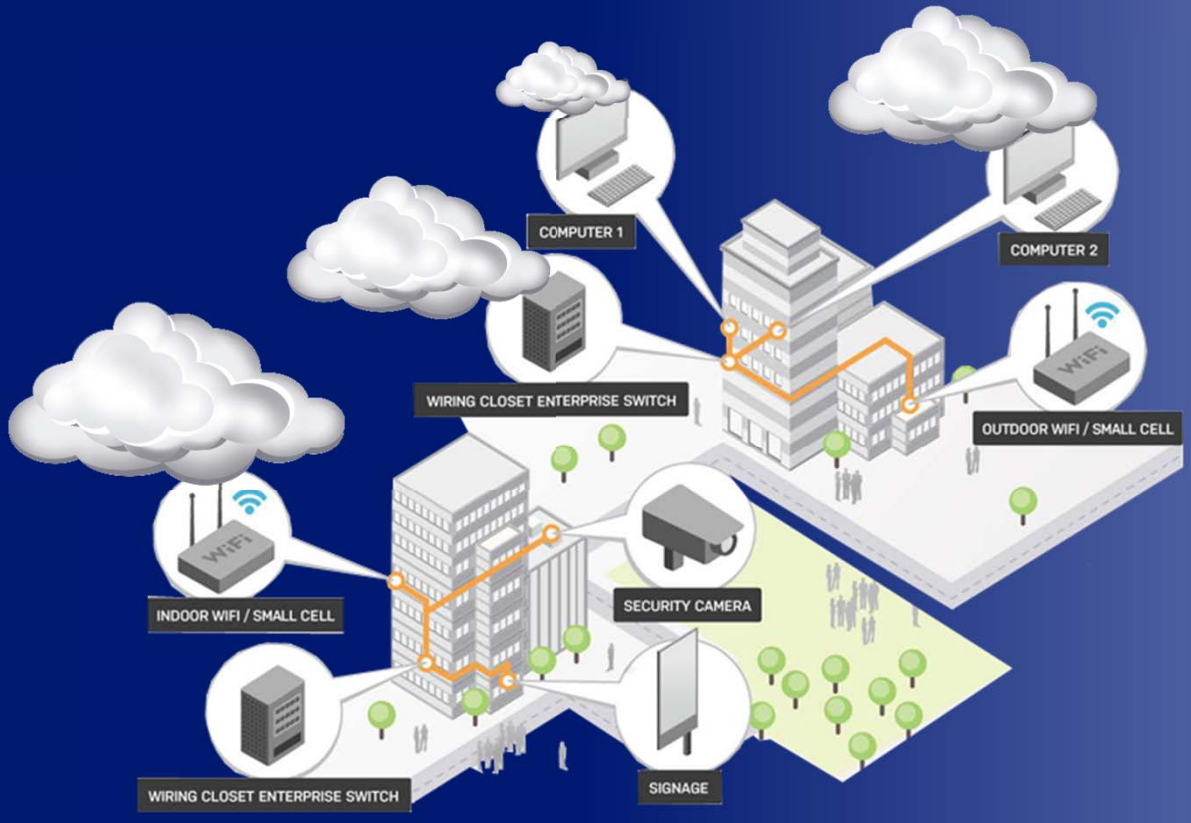

AQUANTIA®

1000BASE-T/2.5G/5.0G/10GBASE-T Signaling

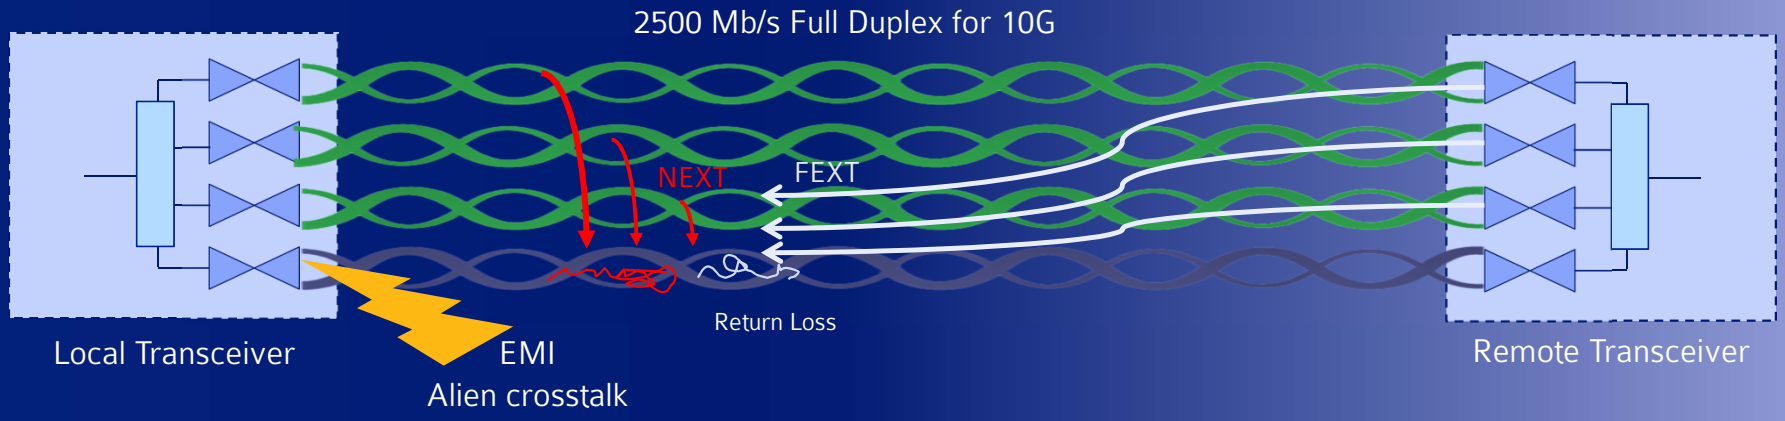

- Duplex transmission → Echo power >> Received power
- Self Near End (NEXT) and Far End (FEXT) Crosstalk
- Alien crosstalk noise from adjacent cables
- Environment specific interferences (e.g. Enterprise)

AQUANTIA®

ESD Impulse Noise in Enterprise

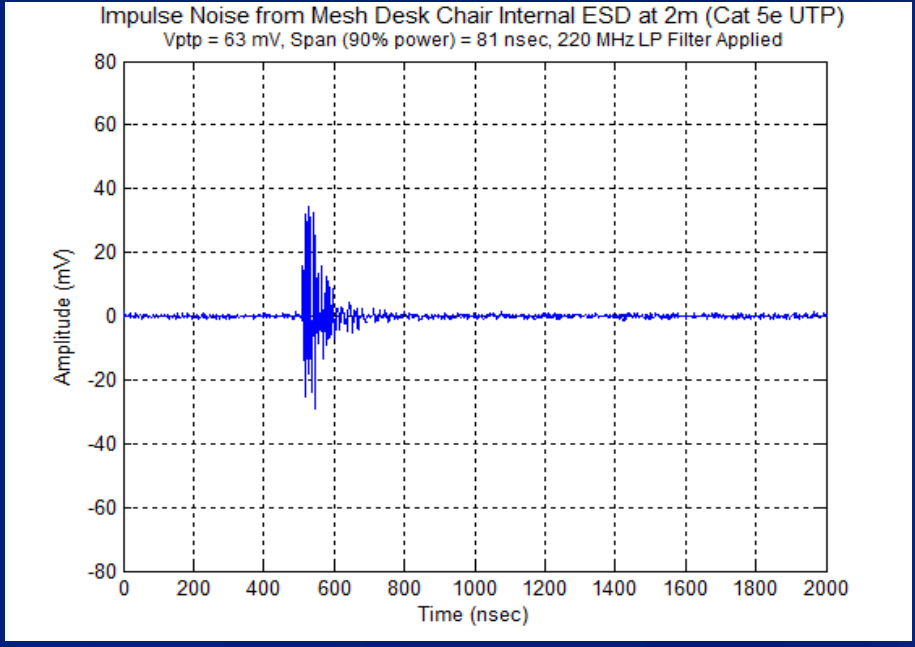

Example of ESD Differential noise in enterprise

- ESD events are generally band pass (80MHz to 200 MHz), low duration, frequent in enterprise that interferes with the operation of data rates above 1000BASE-T
  - Such interference events increases the bit error rate of an otherwise properly operating data link
  - Example enterprise space: Visitors taking a seat & getting up in a public library

AQUANTIA®

Cabling Quality: Data Centers vs. Enterprise

- Next generation Data Centers are built cleanly from scratch with new cabling

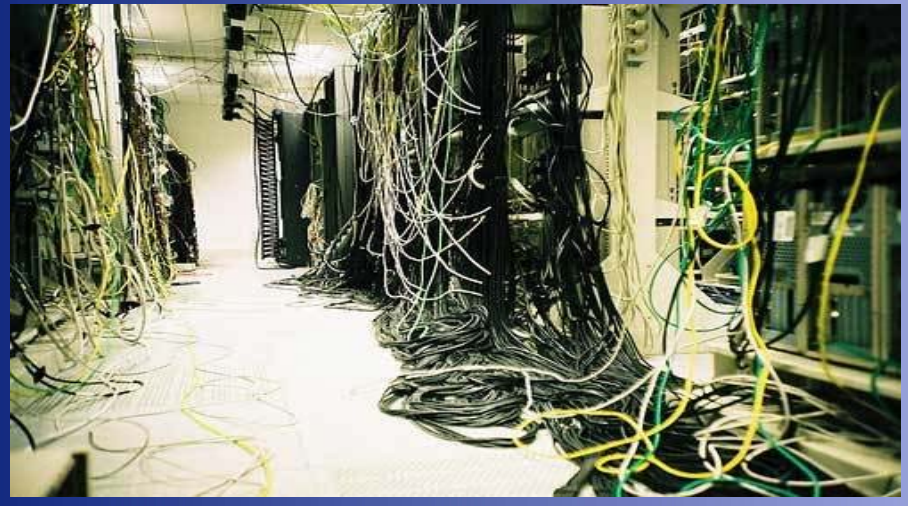

- Enterprise's goal is to extend the life of existing cabling to the next generation
- Many enterprise cabling setups are qualified only up to 1000BASE-T (<62MHz)
- Poor cable return loss leads to large impulse reflections in full-duplex links

AQUANTIA®

Non-linearity in Transmit Signal & Return Loss

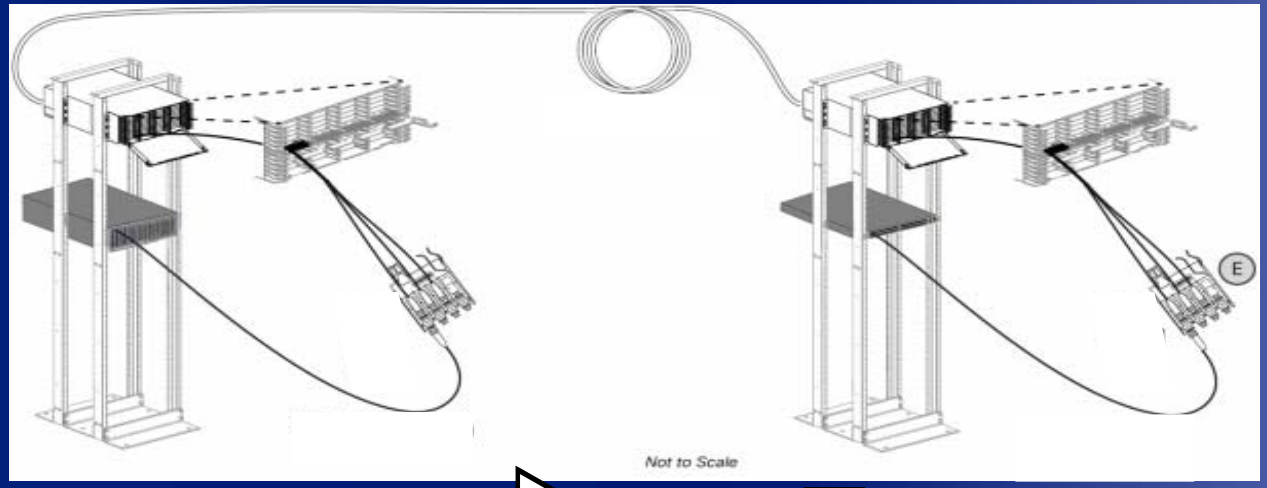

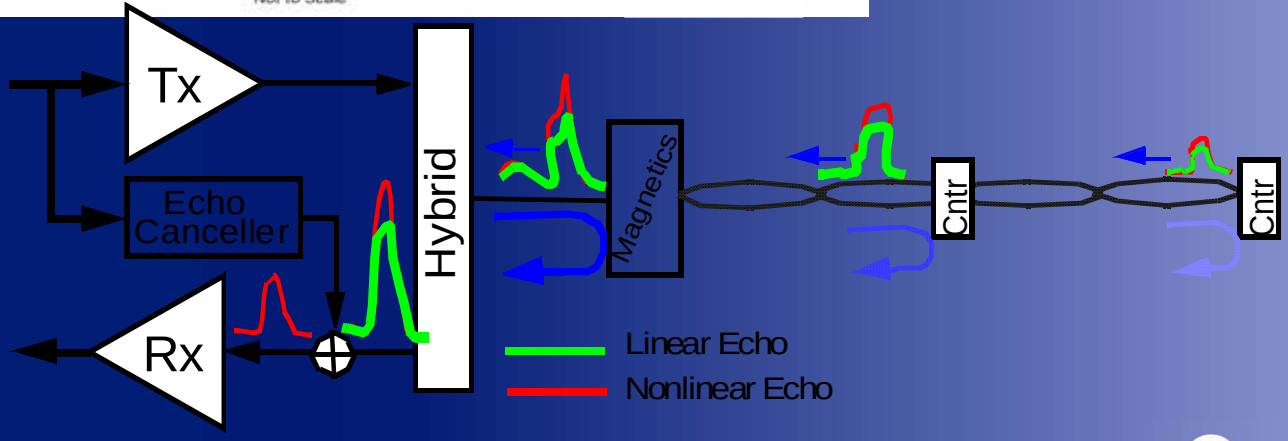

AQUANTIA®

10GBASE-T Coding

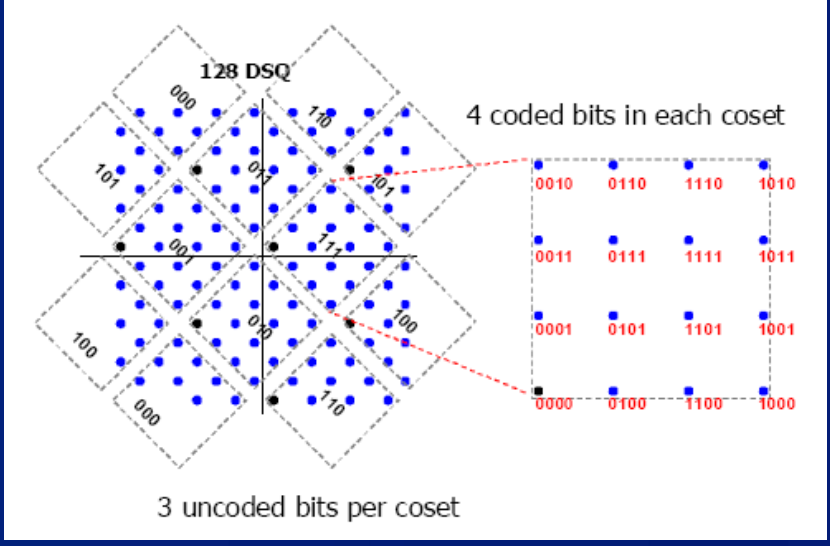

- 10GBase-T DSQ128 allows 128 combinations
- There are 7 bits/DSQ128 symbol
  - 3MSBs are uncoded
  - 4LSBs are FEC coded through LDPC
- These 3 uncoded bits are most vulnerable to large impulse noise dominant in Enterprise

AQUANTIA®

How to Protect All Transmit Bits with LDPC?

- Gray-coded PAM16 signaling
  - PAM 16 = 4 bits per symbol
    - 8bits vs 7bits per 2 symbols
  - Additional bits per frame used to protect uncoded bits by same LDPC machinery
  - Combination of LSB & MSB bits in the LDPC Frame together higher number of encoding bits improves FEC gain by over 1dB
- Robust to non-stationary impulse noise & AFE imperfections
- Otherwise, a scaled version of 10BASE-T

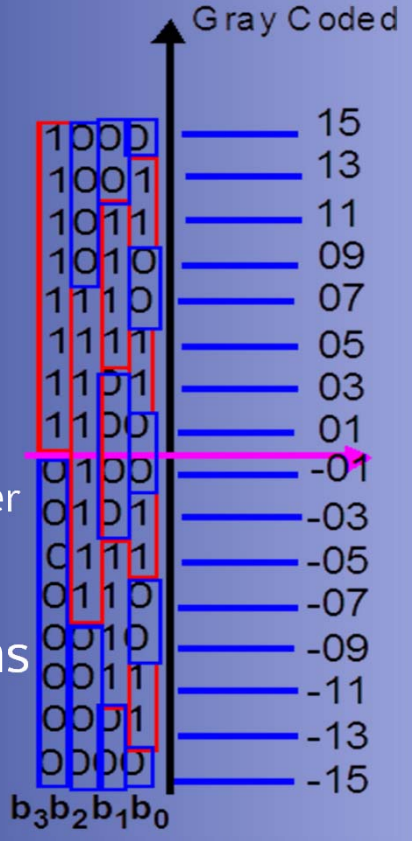

AQUANTIA®

Data Mapping

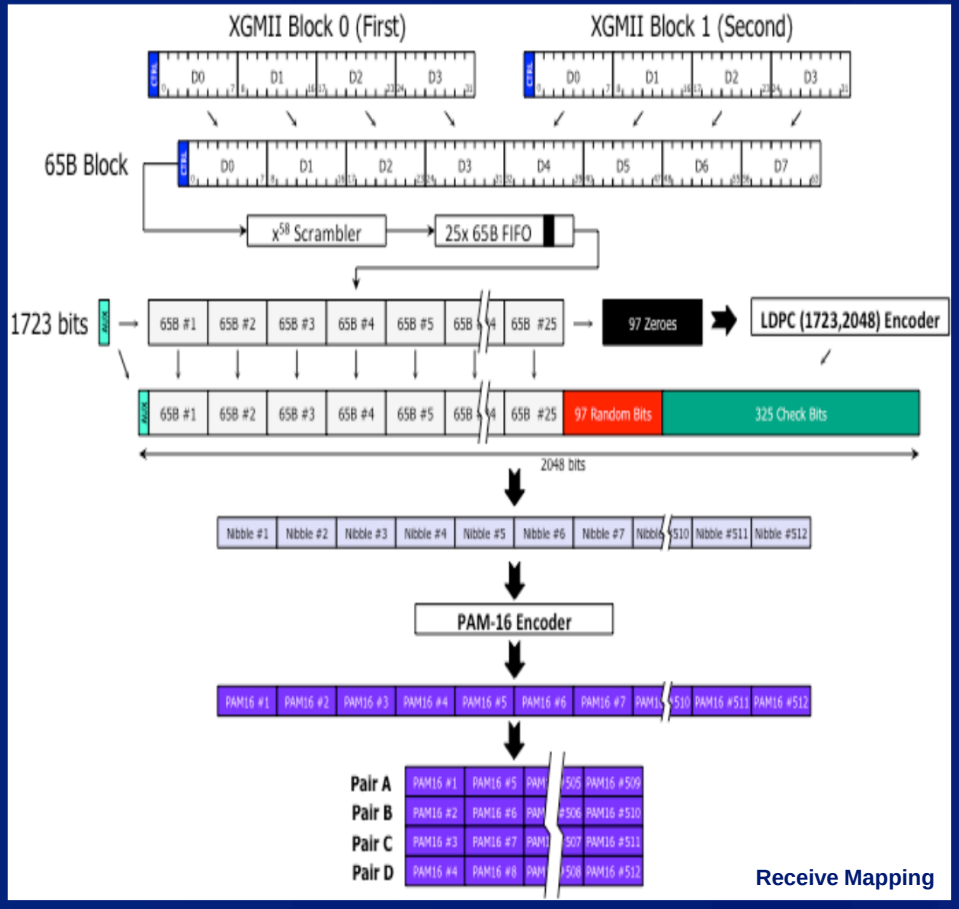

- 5Gb/s via fully LDPC coded PAM 16 running at 400Ms/s
- 2.5Gb/s via fully LDPC coded PAM 16 running at 200Ms/s
- Fully Encoded PAM16 An FEC extended to uncoded bits
  - Removes errors caused by noise spikes
  - Provides additional margin for ESD events
  - Relaxes the PHY transmit linearity spec significantly
  - Relaxes the magnetic RL requirement

AQUANTIA®

Aquantia PHY Block Diagram

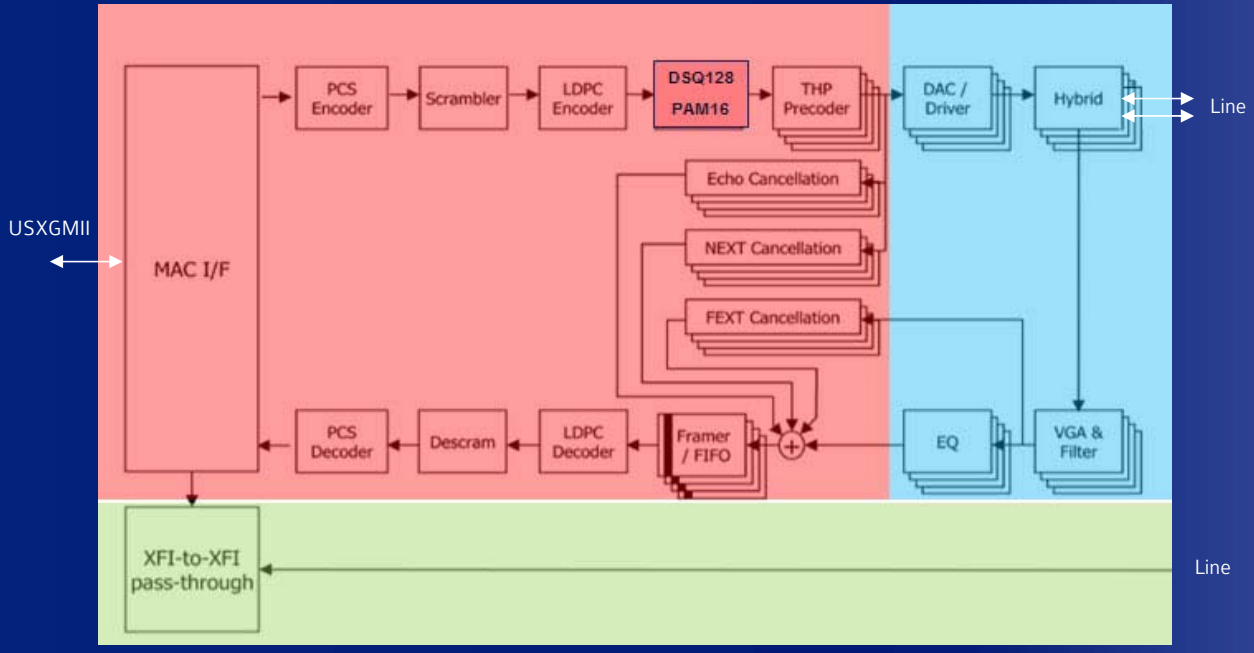

- PHY both analog & digital blocks intelligently architected to power scale proportional to data rates 10G/5.0G/2.5G
  - Analog power saving achieved by modifying performance requirement
- USXGMII SerDes runs at fixed 10Gbps at all PHY data rates to simplify system interface

AQUANTIA®

Aquantia PHY Chip

- Single PHY die view in 28nm process
- Analog frontend with 5 receive channels and 4 transmit channels
  - 5th receive to sense/cancel in band RFI plus on-chip LC filters to cancel out of band RFI
  - Combination of LC & ring PLLs to cover all required speeds
- Strong FEC with 9dB gain (LDPC)
- Modular FIR filters combined with convergence engines
- MACSEC Engine
- Chip Area: 40mm^2
- Power: 1W, 2w, 3W for 2.5G, 5G, 10G
- Supporting 1588 & PTP Features
- Supporting Energy Efficient Ethernet (EEE)

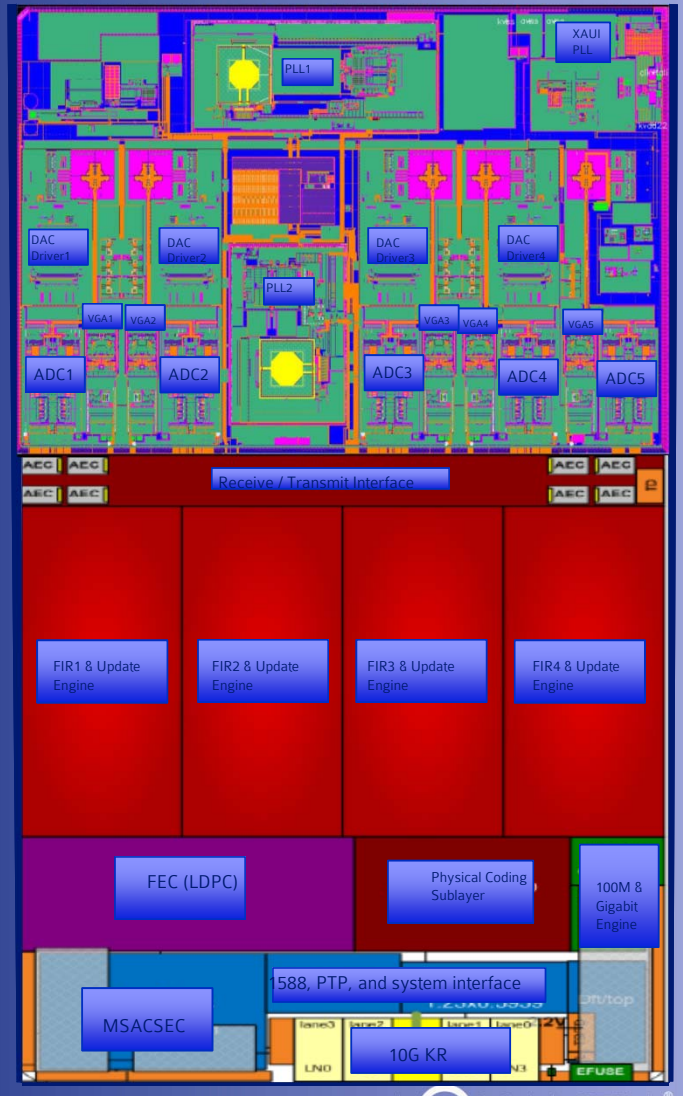

AQUANTIA®

IEEE Performance Specification for 2.5G/5.0G

- 2.5 Gb/s PHY specified for operation over
  - Up to at least 100m on four-pair Class D (CAT5e) balanced copper cabling on defined use cases and deployment configurations

- 5 Gb/s PHY specified for operation over
  - Up to at least 100m on Class E (CAT6) balanced copper cabling on defined use cases and deployment configurations
  - Up to 100m on Class D (CAT5e) balanced copper cabling on defined use cases and deployment configurations

AQUANTIA®

Cat5E & Cat6 Single-Cable Performance Setup

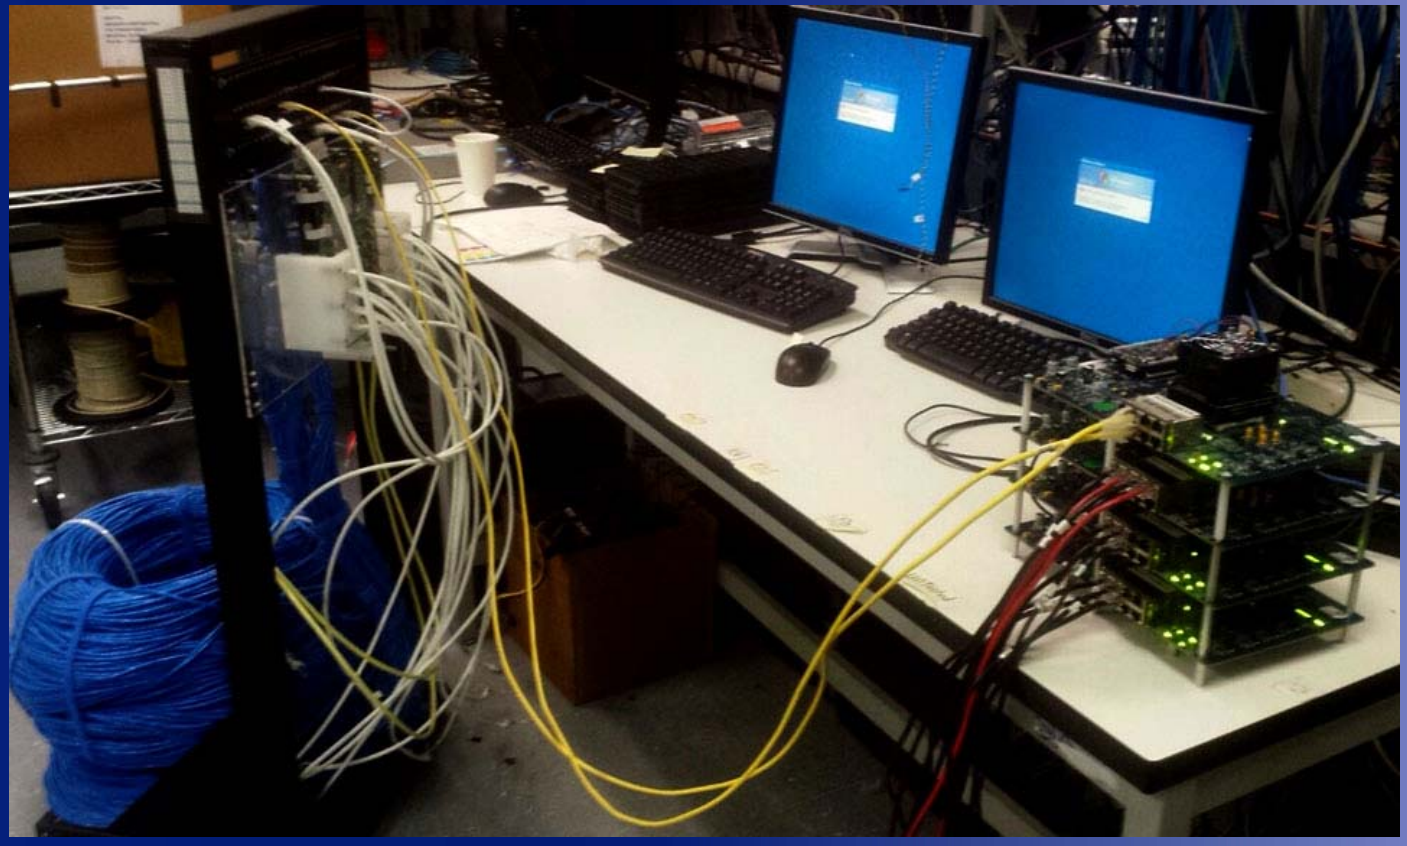

AQUANTIA®

2.5Gbps/5.0Gbps Performance Over Single Cable

- 2.5Gbps reach for BER<1E-12
  - Over ~195m of Cat5E
    - Cable IL = ~42dB@100MHz (18dB below limit line)
- 5.0Gbps reach for BER<1E-12
  - Over ~125m of Cat5E
  - Over ~135m of Cat6
    - Cable IL = ~43dB@250MHz (9dB below limit line)

AQUANTIA®

6around1 Cable Compliance Racks (CAT6A/CAT6/CAT5e)

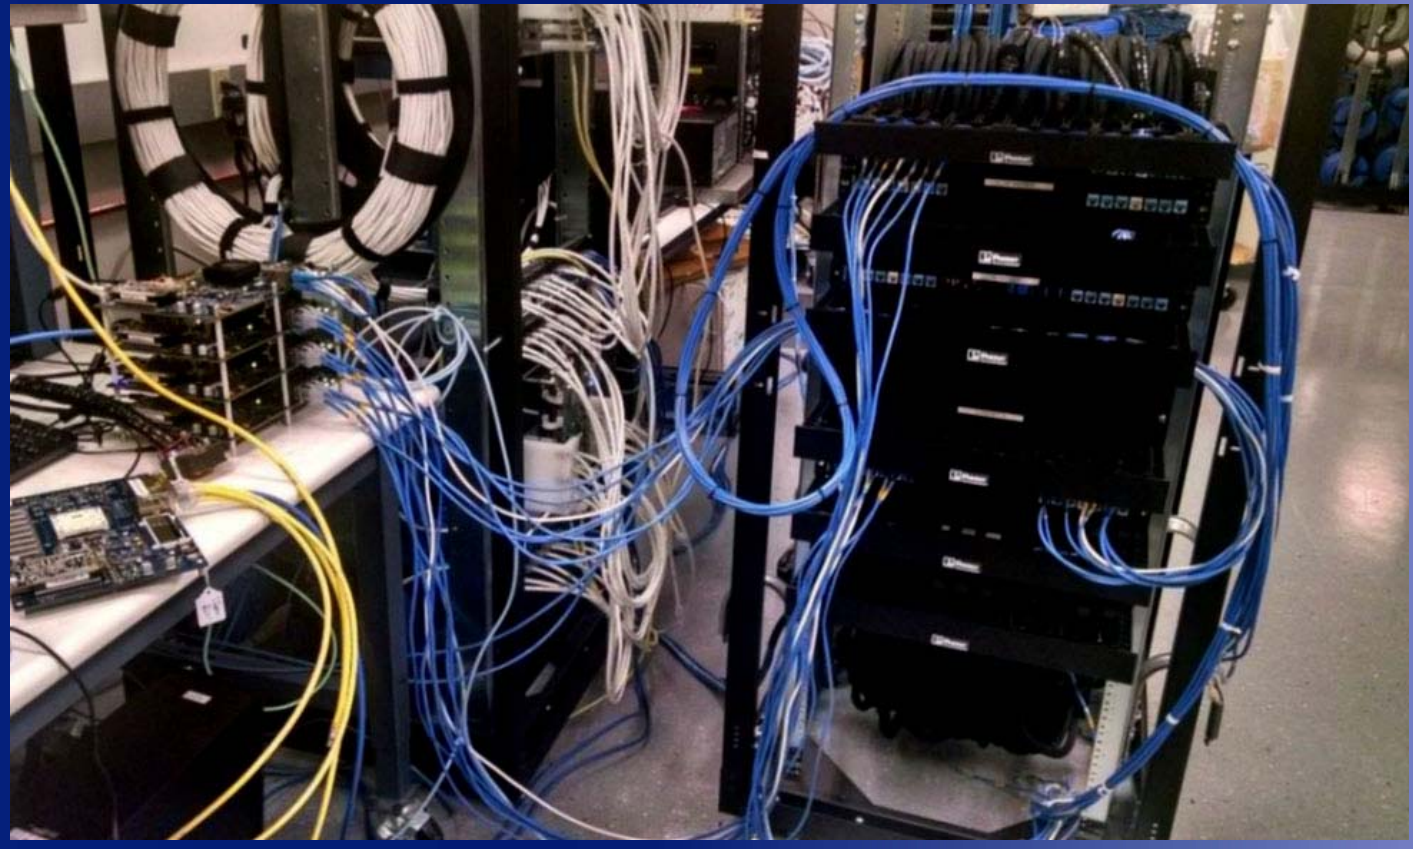

AQUANTIA®

2.5Gbps/5.0Gbps Performance on 6Around1 Cable

- 2.5Gbps reach for BER<1E-12
  - Over ~135m of Cat5E (Full length 6@1)
  - Cable IL = ~29dB@100MHz
    - → Margin to Cat5E Limit= -5dB
- 5.0Gbps reach for BER<1E-12
  - Over ~100m of Cat6 (Full length 6@1)
  - Cable IL = ~32dB@250MHz
    - → Margin to Cat6 Limit= 2dB

AQUANTIA®

2.5Gbps Alien X-Talk Limit Lines (BER<1E-12)

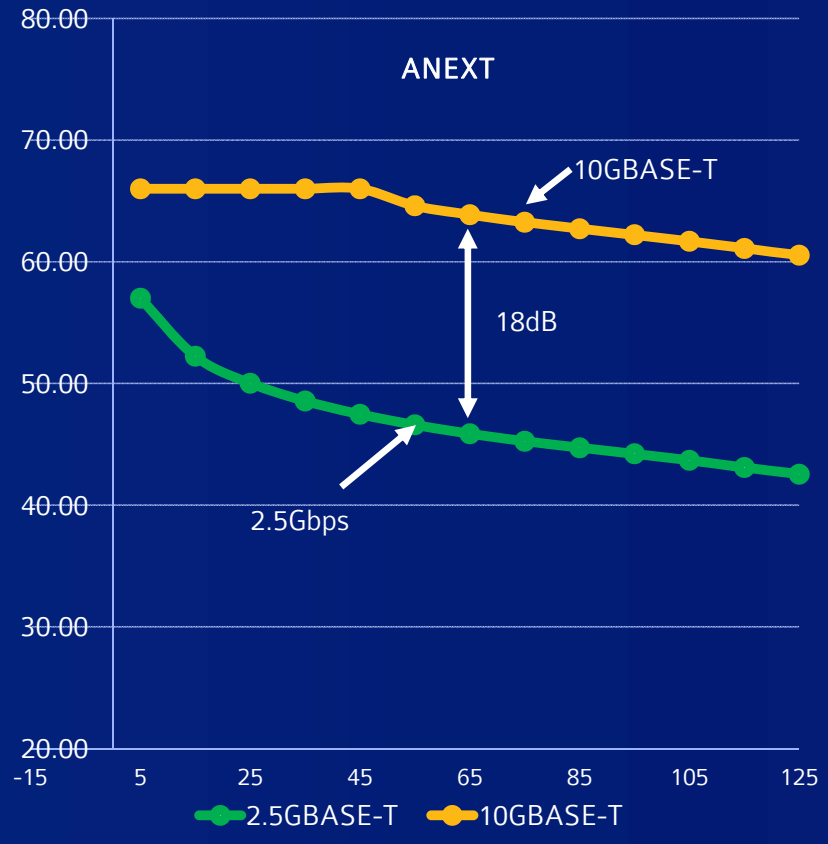

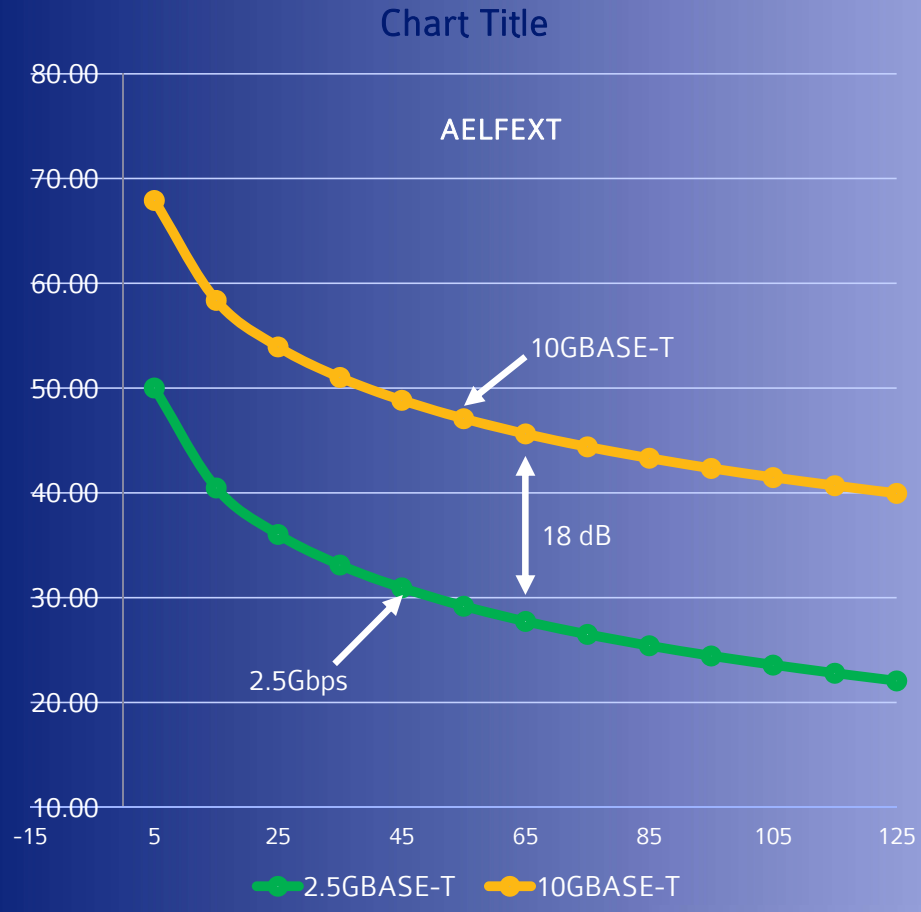

AQUANTIA®

5.0Gbps Alien X-Talk Limit Lines (BER<1E-12)

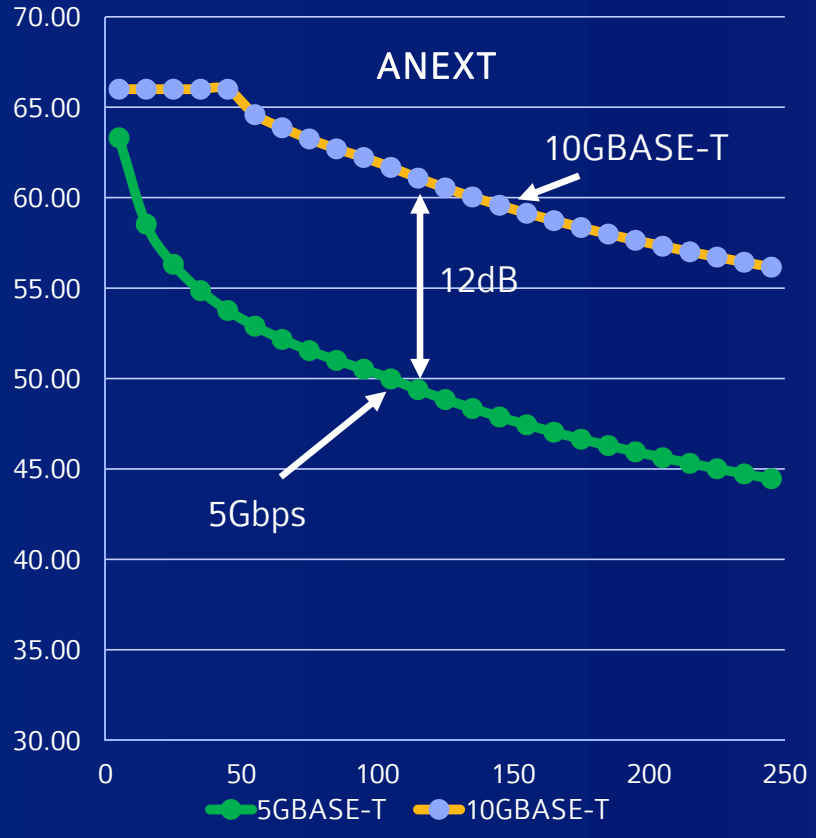

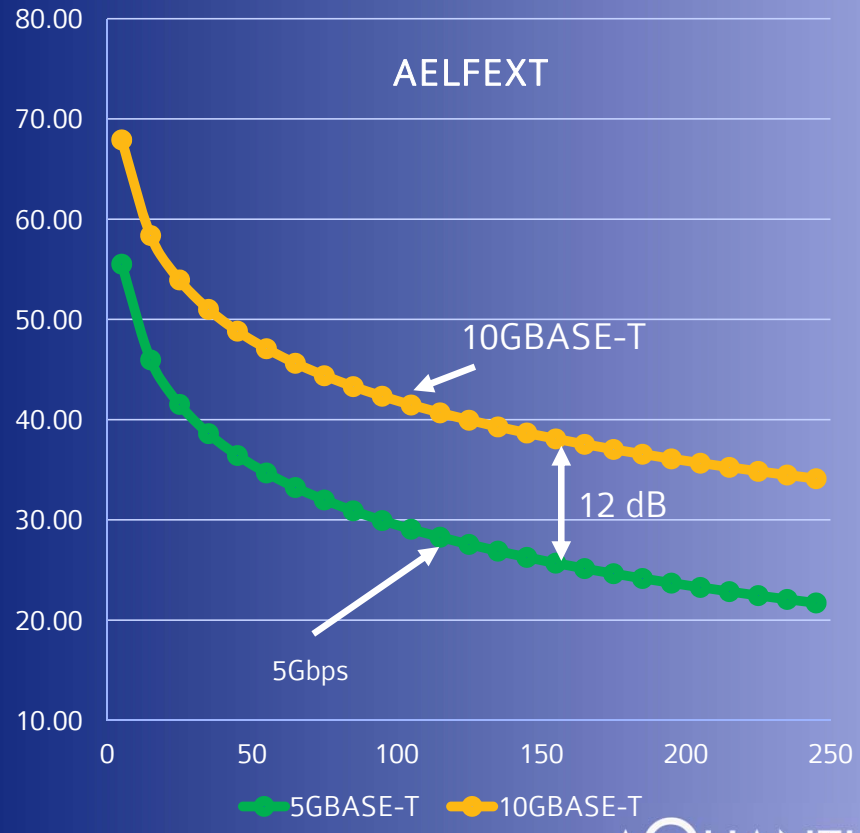

AQUANTIA®

Alien X-Talk of Full 100m 6around1 Cable Setups

ANEXT

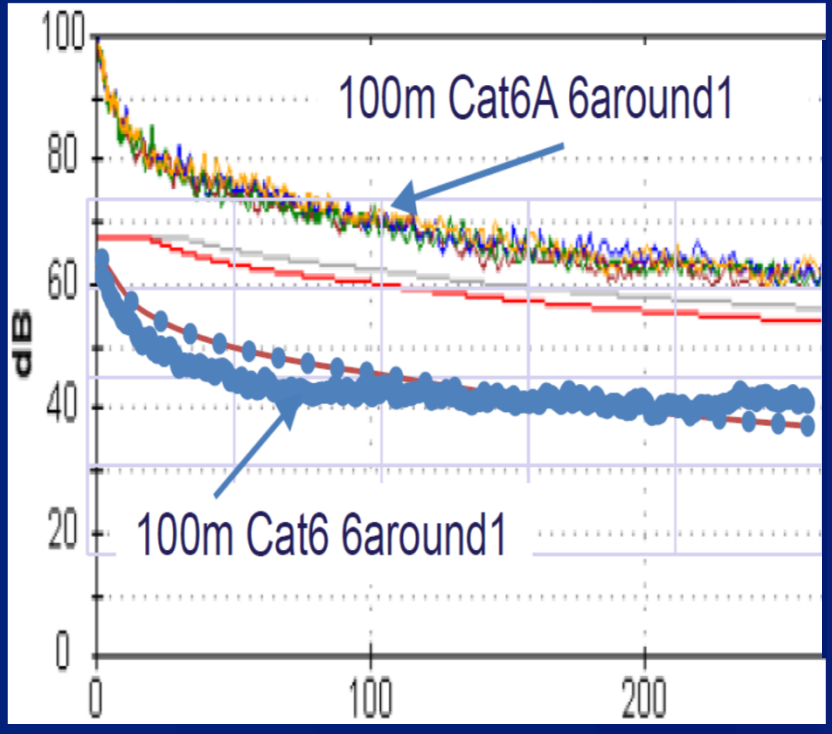

AELFEXT

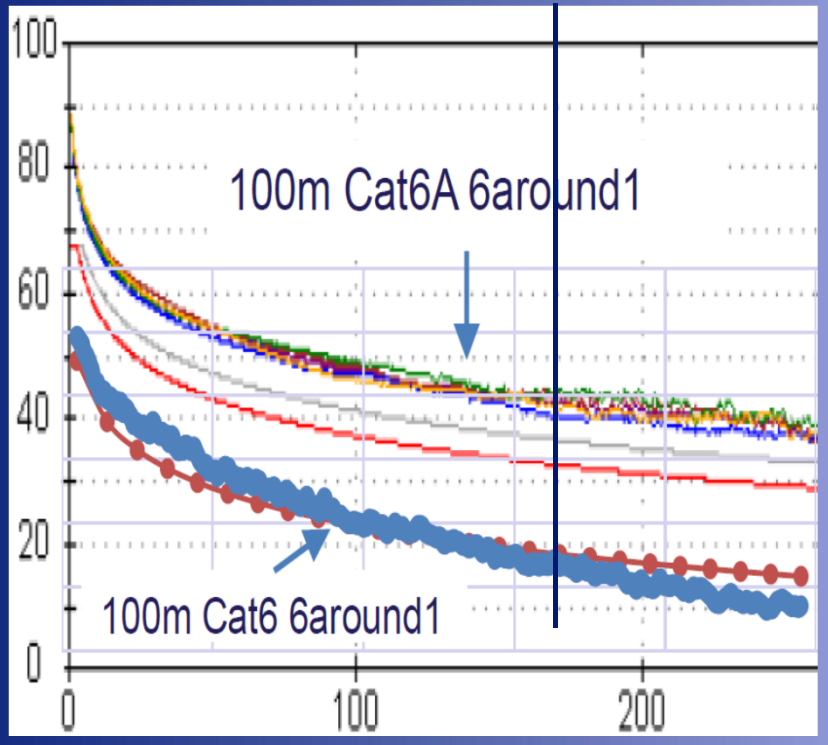

AQUANTIA®

6Around1 Cable Bundled Configurations

- 6around1 bound by two binders all across

- 6@1 tied by straps every 4ft (Common)
  → ~9dB improvement in alien xtalk

- unbound cables as in conduits
  → >20dBimprovement in alien xtalk

http://www.ieee802.org/3/10GBT/public/mar03/vanderlaan_1_0303.pdf

AQUANTIA®

IEEE Defined WAP Enterprise Use Case

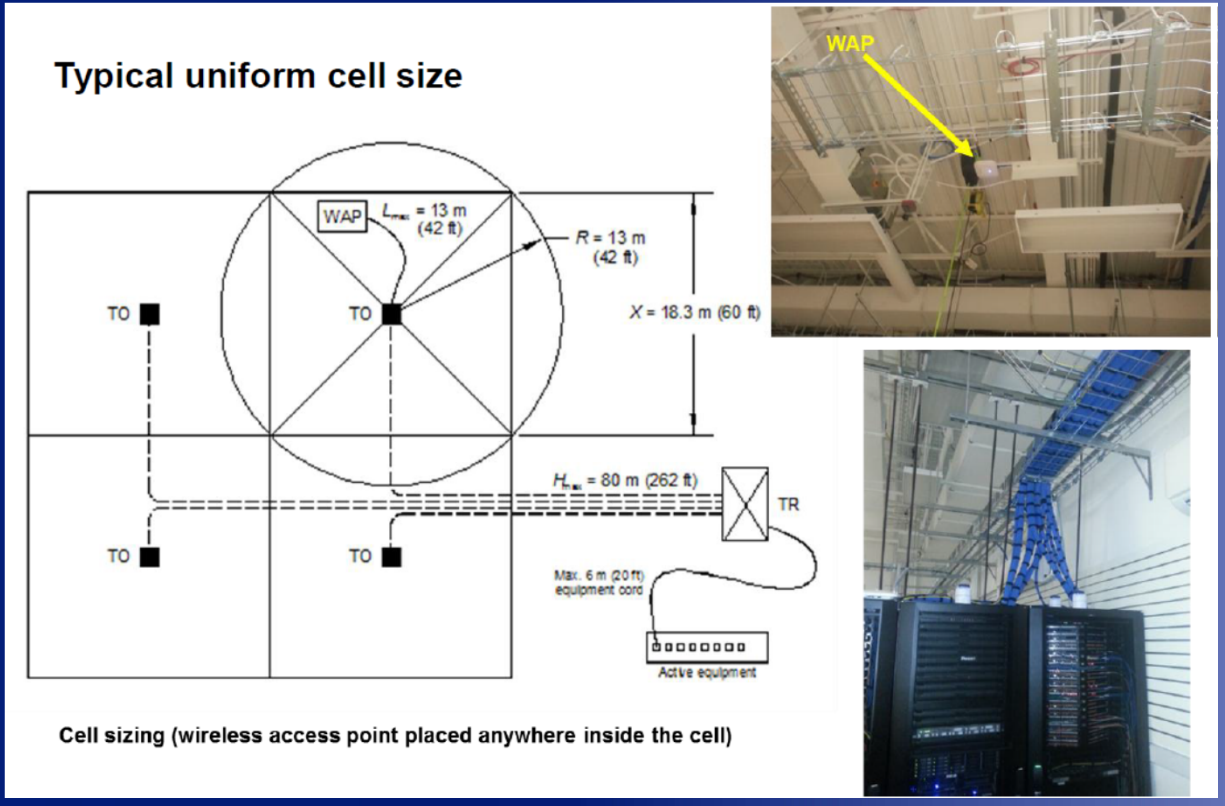

http://www.ieee802.org/3/NGBASET/email/pdfy0VqAHwlWI.pdf

AQUANTIA®

CAT6 Alien X-Talk: Typical Enterprise Installations

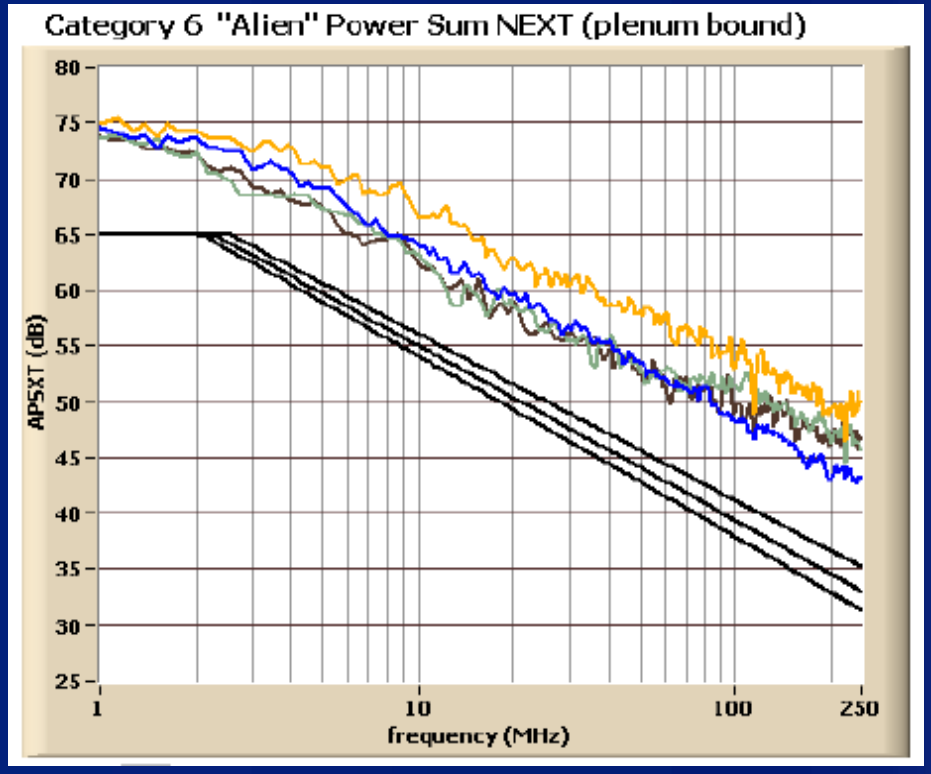

- Configuration:
  - 24 cat6 Cables
  - Tied every 5 feet
  - Tied length <80m
  - Victim at center of bundle
  - Routed to the ceiling
- X-Talk Measurement meets the target 5Gbps alien xtalk limit lines!

AQUANTIA®

Conclusion

- The ever growing data requirements in the cloud and mobility is creating a bandwidth constraint in global IT infrastructure
- In 2012, Aquantia focused on adding 2.5G/5G speeds on cat5e / cat6 Links to address the next generation wireless access bottleneck
- Aquantia's five speed ICs (single, dual, and quad) are the only commercially available NBASE-T compliant ICs since 2013
- Five speed ICs from Aquantia has been shipping in large volumes in all major switch and AP platforms since 2014